INTEL'S REVIEW OF IBM'S CLAIMS[1]

I. Overview

Intel has continued its review of IBM's claims that contradict Intel's detailed analysis of the Pentium™ processor floating point flaw. Actual data from real world applications confirms Intel's estimates, and contradicts IBM's claims.

IBM questions Intel's estimates in two ways: (i) IBM asserts that certain bit patterns susceptible to the Pentium processor's floating point divide flaw occur with increased frequency in the course of calculations such as those typically performed in a spreadsheet application with financial data; and (ii) IBM also asserts that typical spreadsheet users perform divide calculations at a dramatically higher rate than estimated by Intel.

For this review we studied data drawn from real spreadsheet applications in representative user environments (this is in addition to the exhaustive analysis described in the Intel white paper). Data has been collected from hundreds of spreadsheets throughout Intel. Our sample applications were collected from the various functions within Intel, including Finance, Sales & Marketing, Planning, Treasury, Product Engineering, Production Control and Tax/Customs.

II. Frequency of Suspect Bit Patterns in Denominator

Bit patterns susceptible to the floating point divide flaw all contain a string of more than six 1's in their binary representation (although not all numbers with this pattern are at risk). To test the IBM assertion that denominators with suspect bit patterns occur more frequently than random, we measured the bit patterns fed as denominators to all the floating point divides encountered in fully recalculating our sample of spreadsheets. Our findings are presented in Figure 1.

Figure 1

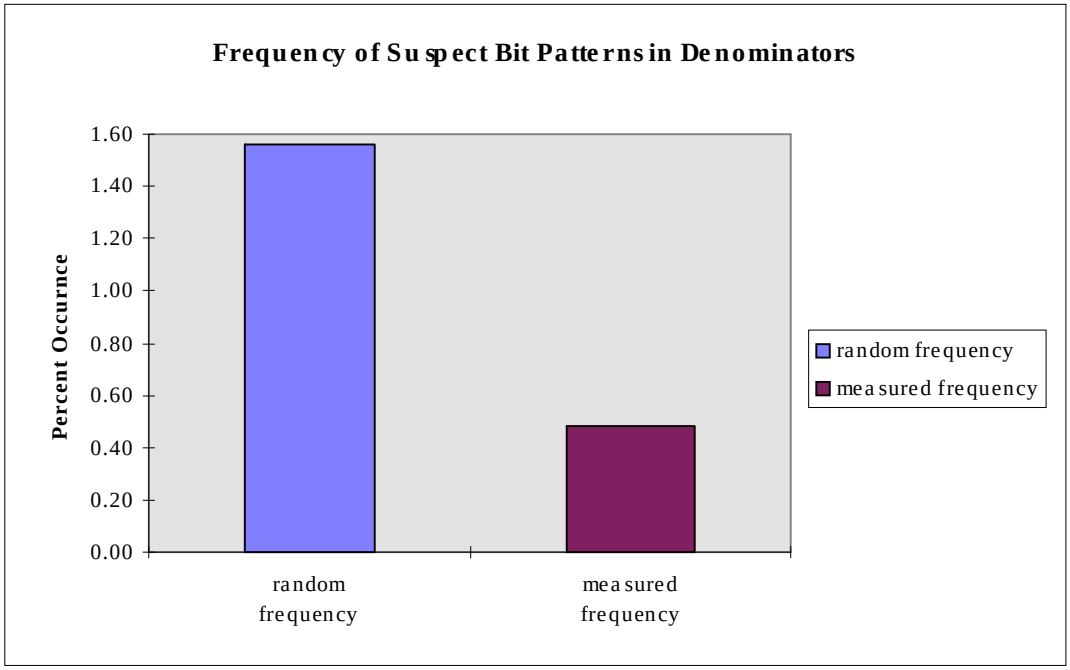

[1] This is a summary of a paper entitled "Analysis of Floating Point Divide Calculations in Spreadsheet Applications in the Commercial PC Marketplace" by M.L.Barton, Ph.D., Staff Computational Scientist, and R.A.Passov, Senior Treasury Manager, Intel Corporation.

As this figure shows, the measured frequency of the bit patterns that are necessary (but not sufficient) for an at risk denominator is 0.48%, actually below the expected random frequency of 1.5625%. Put another way, while IBM asserts that the frequency of encountering floating point divide errors based on denominators with several consecutive 1's in spreadsheets is 1 in 100 million vs. Intel's claim of 1 in 9 billion, the measured frequency is in fact even less than 1 in 9 billion.

In their analysis IBM focused on expressions such as (a.b) - (c.d), an example of which is 4.1-1.1, which is calculated to be 2.9999999999999. This number has a long string of 1's in its binary representation. However, not only does the actual data show that these numbers are no more likely to occur than other numbers as a denominator, but in fact when such a number does trigger a reduced precision result, the inaccuracy shows up at worst in the 12th significant decimal digit. IBM neglected to mention this fact.

III. Frequency of Floating Point Divide Operations

A second IBM assertion is that a spreadsheet recalculation running on a Pentium processor-based PC can perform about 5000 floating point divide operations per second. Based on this, IBM assumes that a typical user will execute recalc's for 15 minutes/day, thus performing 4.2 million floating point divide operations daily. IBM presents no data to support this assumption. Intel has concluded that a basic spreadsheet user will perform approximately 1000 floating point divides per day.

As the measurements in Figure 2 indicate, nearly 70% of the spreadsheets in our survey involved fewer than 1000 floating point divide operations, and virtually all of the spreadsheets (other than the most intensive financial and engineering applications) involved fewer than 10,000 divides.

Figure 2

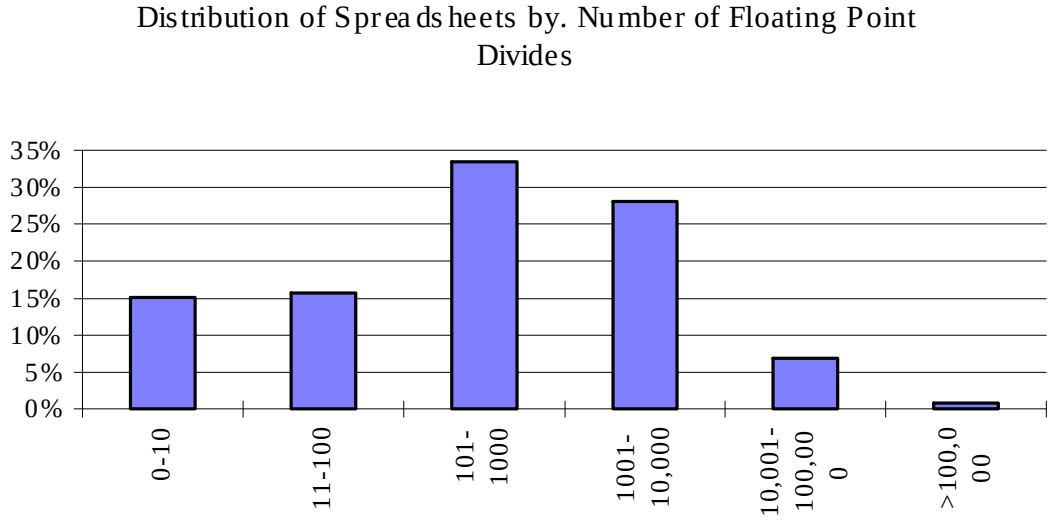

Based on these measurements and our ongoing survey of how users interact with these applications in terms of the number of spreadsheets opened per day by a typical user (approximately 1-3), the number of recalc's performed per spreadsheet (approximately 1-3), and the percentage of overall formulae involved in a recalc (about 25%), we remain confident of our conclusion that the basic spreadsheet user performs fewer than 1000 floating point divides in a day.

IV. Conclusion

IBM made two assertions counter to Intel's white paper that led to IBM's conclusion that a computer based on the Pentium processor could return a reduced precision result every 24 days. The data presented here, which supplements the data provided in the Intel white paper, proves that IBM's assertions are incorrect, and strongly supports Intel's analysis of the impact of the Pentium processor's floating point divide flaw. This flaw is of no significance in the commercial PC market as the average spreadsheet user is likely to encounter a failure only once every 27,000 years.